

SCALABLE DISTRIBUTED DATABASE SYSTEM AND METHOD FOR LINKING CODES TO INTERNET INFORMATION

[0001] This application claims priority to U.S. Provisional Patent Application No. 60/252,906 entitled "Scalable distributed database system for linking URNs to Internet information" filed November 27, 2000, which is hereby incorporated by reference.

FIELD OF THE INVENTION

[0002] The present invention relates to connecting an end user with locations on the Internet based on codes inputted by the end user.

BACKGROUND OF THE INVENTION

[0003] Bar code scanning technology has been in use for quite some time. Such bar codes appear on most items encountered by consumers every day – from grocery items to driver's licenses. Other types of codes (such as private codes), and methods for acquiring such codes, have been and are continuing to be developed. Bar codes and other codes allow for fast and easy identification, tracking, and inventory of items as well as storage of data relating to such items. As bar codes have become a common and familiar part of the consumer experience, so has the use of web services and applications. Different types of web services and applications continue to grow in number and popularity.

[0004] Identifying ways of converging these two apparently disparate technologies to create a product that provides mobility, simplicity, accuracy, convenience and efficiency to web applications and services would enable companies to reduce costs (by replacing formerly manual tasks with bar coding technology), increase revenue (by offering barcode-enabled advertising alternatives that allow consumers to learn about products, prices, store locations etc.), increase productivity (by automating sales force tasks such as product ordering), and improve customer loyalty (by offering barcode related value-added services to customers).

SUMMARY OF THE INVENTION

[0005] The present invention provides systems and methods related to the convergence of bar code and other coding technologies and web-based applications and services, thereby fulfilling a need present in the prior art.

[0006] In particular, the present invention is directed to a method and system for directing an end user to a location on a network based on a code inputted by the end user. The end user is associated with at least one of the following: an end-user country selection, an end-user language selection, a

service category, an access media type, and the presence of an access certificate. The access media can be one of at least two types (at least one of the types being a bar code scanner). A database that includes a plurality of entries is provided. Each of the entries is associated with a machine-readable code and includes one or more different network addresses. The different network addresses are associated with different providers. Code information associated with one or more scans of machine readable codes, made by the end-user with a remote scanning device, is received at a network site. In response to the scan information, entries associated with the machine-readable codes scanned by the end-user are retrieved from the database. At least some of the entries include information that is not relevant to at least one of the end-user country selection, the end-user language selection, the service category, the access media type, and the presence of an access certificate. Based on at least one of the end-user country selection, the end-user language selection, the service category, the access media type, and the presence of an access certificate, one or more network addresses are selected from the retrieved entries. The end user is connected with the one or more network addresses selected.

[0007] It is to be understood that both the foregoing general description and the following detailed description are exemplary and explanatory and are intended to provide further explanation of the invention as claimed.

BRIEF DESCRIPTION OF THE DRAWINGS

[0008] The accompanying drawings, which are included to provide further understanding of the invention and are incorporated in and constitute a part of this specification, illustrate embodiments of the invention and, together with the description, serve to explain the principles of the invention.

[0009] In the drawings:

[0010] Figure 1 illustrates the manner in which links may be organized in accordance with a preferred embodiment of the present invention;

[0011] Figure 2 illustrates a preferred embodiment of a link that can be used in connection with a preferred embodiment of the present invention;

[0012] Figure 3 illustrates the manner in which a user may be precluded from obtaining access to certain services in accordance with a preferred embodiment of the present invention;

[0013] Figures 4A – 4C illustrate an exemplary data structure that may be used to design a database in accordance with a preferred embodiment of a system of the present invention;

[0014] Figure 5 illustrates a cache database schema that may be used to design a database in accordance with a preferred embodiment of a system of the present invention;

[0015] Figure 6 illustrates a preferred embodiment of a system of the present invention; and

[0016] Figure 7 illustrates a flow chart of a preferred embodiment of a method for connecting end users to Internet information based on codes inputted by the end users.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

[0017] Reference will now be made in detail to the preferred embodiments of the present invention, examples of which are illustrated in the accompanying drawings. Wherever possible, the same reference numbers will be used throughout the drawings to refer to the same or like parts.

[0018] The present invention is directed generally to connecting a unique identifier, such as a bar code, to a URL that points to a web page on the Internet. The connection is tailored based on a number of different criteria including, for example, a particular language of the end user, the country of the end user, service categories chosen by the end user, and the media by which the end user accesses the system.

[0019] Figure 1 is a diagram illustrating a preferred embodiment of the manner in which URLs may be organized within the inventive system. Other categorizations, groupings and organizations will be known to those skilled in the art and are within the scope of the present invention. Link layer 101 represents a collection of one or more unique identifiers linked to one or more URLs (the combination of one or more unique identifiers linked to one or more URLs is referred to herein as a "Link"). In the preferred embodiment, the Links are grouped into different languages. Figure 2 illustrates an example of a Link in which three EAN codes 201, 202 and 203 (i.e. unique identifiers) are linked to two different URLs 204 and 205 (one in English and one in French). In one embodiment, the unique identifier consists of a namespace and a number string that is namespace unique. This allows the inventive system to store in a database many types of unique identifiers (e.g. for barcodes, RFID tags, credit card numbers etc.). For example, an EAN-barcode, number 7331040021186, is represented by the namespace EAN13, and the string 7331040021186. By way of further example, a credit card issued by Company ABC with the card number 1234 5678 6543 would be represented by the namespace ABC and the string 123456786543. Other types of unique identifiers can be used in accordance with the present invention.

[0020] Service layer 102 contains collections of Links that are specific to particular companies (i.e., companies sponsoring the web pages to which the URLs point), each sharing properties of a particular service (i.e., groupings of Links that share a common characteristic). The shared characteristic may be, by way of example, an end user country of origin, language preference, profession, or access media device type. Other types of services based on any shared characteristics can be defined in accordance with the present invention.

[0021] For example, a French medical company marketing a painkiller product available only in France may create a service pursuant to which URLs pointing to marketing information on its painkiller product are presented only to users in France. Thus, the medical company could, for example, define two Links – a global Link and a French Link. A French user would be directed to URLs within the French Link, but everyone else in the world would be directed to URLs within the global Link. The medical company would, thus, have the ability to engage in a country-wide campaign for the painkiller for France, but still maintain the information about the product on the Internet for the rest of the world. A Canadian company may want to create two language preferences services, namely, a service for Links containing URLs directed to web pages written in English and a service for Links containing URLs directed to web pages written in French. By way of further example, a medical company may wish to present to doctors different URLs pointing to information for a painkiller product than those it presents to patients. The medical company would then create two services (one for doctors and one for patients) each of which holds a Link for different information relating to the painkiller product.

[0022] In some embodiments, a company may want to create services for Links to which only certain users or groups of users can obtain access. This can be accomplished using certificates. For example, one service may be created for Links containing URLs that can be accessed only by individuals that present a particular certificate. In a particular example, the French medical company may create a service that includes Links containing URLs that can be accessed only by individuals with a certificate proving that the individual is a doctor. Figure 3 illustrates the manner in which a patient 301 has access to buy Link 302 and medical Link 303. However, doctor 305 has access not only to buy Link 302 and medical Link 303, but to doctor Link 304, because doctor 305 has certificate 306. Another service may be created that includes Links containing URLs that can be accessed only by end users using a particular device. For example, each barcode reader contains a serial number that is transmitted together with the barcode. Thus, certain Links may be protected by a service that provides access only to end users employing readers with certain serial numbers. This embodiment may be advantageous to a company seeking to make sure only its readers can be used for obtaining access to information about its products.

[0023] Category layer 103 controls how URLs are presented to the end user. Each service in service layer 102 may belong to a particular category. A category is a group of services that all share a common characteristic. The end user may select the categories of Links with which the end user is presented. Thus, for example, the end user may indicate that it is only interested in being presented with URLs relating to a particular product or, for example, to shopping sites.

[0024] One exemplary category is a product category 105 (shown in Figure 1). All services relating to a particular kind of product will be grouped into product category 105. Information category 106 groups services that pertain to third party information about the service. Thus, for example, upon scanning a barcode for a video tape of a movie and choosing an information category, the end user will be presented with URLs pointing to reviews of the movie. Shopping category 107 groups services pertaining to Links with URLs identifying where an item can be purchased. Another type of category, an application category, groups Links by application, rather than merely by service, which implies a more advanced interface to the system. The API of the present invention provides access to a number of system-, device- and user parameters. A call to an application can be of several types, a simple re-direct or a general call with a flexible protocol. The APIs may differ depending on the type of access. Categories can be added, removed or changed on demand. Categories in addition to those described here can be created and used in the inventive system within the scope of the present invention.

[0025] Access layer 104 controls the URLs that are presented to the end user depending on the access media used by the end user. The different types of access media a user may employ include, by way of example, bar code readers (portable or tethered to a personal computer), personal computers, WAP enabled mobile phones, and personal digital assistants. Each grouping of Links in the category layer are specified as belonging to a certain access type. Thus, Links in a given category are tailor-made for a specific access type, which ensures that each device receives information that the specific device can handle. For example, categories may include Links that are tailored to be shown on the web or using WAP services. In another example, categories may include Links that are tailored to be heard as voice/music using a mobile phone. In particular, a record company wanting to promote one of their artists can have a Link (including URLs pointing to a web site where the artist's music track can be heard) assigned to a category for mobile phones. This music track can then be accessed using a barcode reader connected to a mobile phone. In another example, some categories may contain Links that are tailored to be sent back using Short Message Service ("SMS") to a mobile phone. In particular, an end user may scan a barcode on a book and receive an SMS back with a short review of the book. The end user may scan a barcode on a poster for a concert and receive an SMS back with an electronic ticket for the concert. Other access types can be provided for in the inventive system within the scope of the present invention.

[0026] Preferred embodiments of a central relational database and a cache database that may be used in connection with connecting end users to URLs in accordance with the present invention are described with reference to Figures 4A through 4C and 5, respectively. Figures 4A through 4C

illustrate an exemplary data structure of central relational database 601 (shown in Figure 6) for storing data used in connection with the present invention. Figure 5 illustrates a preferred embodiment of a schema of cache database 603 (also shown in Figure 6). Cache database 603 is designed, in the preferred embodiment, to handle lookup requests with a high degree of performance in terms of speed and capacity. This is accomplished through a combination of the database architecture and RAM caching. In particular, the cache database is, in the preferred embodiment, an object oriented hash based database, not a relational database. Each table in the cache database has a RAM cache.

[0027] With reference to Figures 4A through 4C, central relational database 601 includes a number of tables. Tables 401, 402, 403, 404 and 405 allow companies to set up and access accounts, and create and manipulate data relating to Links. For example, table 401 is used to store data on the company, such as identification information, password, address and other contact information. Each company has associated with it one or more users that enter information into the system. Tables 402 and 404 are used to store preferences for each user in the system (e.g., whether the user prefers help screens, whether the user prefers to have URLs confirmed prior to their entry into the system etc.). Tables 403 and 405 are used to store information relating to a language preference for the user as well as to set templates for the language fields used by the user when entering Link information. Tables 406, 407 and 408 are used to store and manipulate codes implemented by a company in bulk.

[0028] Table 410 is used to store product information and, in particular, Links based on unique identifiers from table 409 and URLs from table 421. For each product identified in table 410, information regarding the company sponsoring the product can be identified from table 411, along with information regarding the user that inserted the product into the system (identified from table 401). Each product identified in product table 410 is associated with a particular service (identified in table 413), and a category (identified in table 412).

[0029] Company data stored in table 411 is associated with data relating to a particular market, stored in tables 415, 416 and 417. Service data stored in table 413 is associated with channel information stored in table 418 and channel type information (i.e. access media device type) stored in table 419, as well as service information stored in table 414. Keyword data, stored in table 420, is used to further separate services that share a common characteristic within a category. For example, two shopping services A (a book store) and B (a pet shop) are separated by the keywords bookstore for service A and pet shop for service B. Service information stored in table 414 relates services of a particular type (defined in table 413) that are owned and shared by a number of companies (defined in table 411).

[0030] The lookup procedure can be described with reference to Figure 5, which illustrates several exemplary tables of cache database 603. Based on a particular code inputted by an end user (e.g., scanned) and the symbology of that code, object dictionary table 501 is used to obtain all entries (each of which includes a service ID and an object ID) from entry table 502. All data associated with the service ID and object ID for each entry is then obtained. In particular, for each object ID, a link is identified from object table 503 and, for each link, link data is obtained from link table 504. For each service ID, a service parent ID is determined from service table 505. The service parent ID allows the inventive system to identify whether a particular service is globally valid or whether the service pertains only to a particular country.

[0031] Also, for each service ID, using service table 505, a company ID (identifying the company that has inserted the service into the system) is identified. Using the company ID, company information is obtained from company table 506. The company information includes, for example, whether a parent company exists as well as whether the company is active in the system. Based on the reference country ID in company table 506, country data is obtained from reference country table 507.

[0032] Returning again to service table 505, a service data ID is identified. Using the service data ID, service data information is obtained from service data table 508. The service data information includes, for example, the name for the service, whether a priority has been set, the type of service, whether any certificates are required to access the service, and a description of the service. The service information (in particular, "pre_url" and "post_url") allows for the construction of URLs from strings defined in the service. "Service hook" determines whether such URLs should be constructed and what type of URL compositing should occur. From service data table 508, a channel ID is determined, which allows channel data (i.e. category information described with reference to Figure 1) to be obtained from channel table 509. From channel table 509, a channel type ID is determined which allows channel type data (i.e. access media device type) to be obtained from channel type table 510.

[0033] The superset of data identified is then filtered to obtain Links that are particular to the current end user session. In particular, as discussed in more detail above, Links that are not associated with the country from which the end user has inputted the code; the preferred language of the end user; the channel (i.e. categories); and/or the channel type (access media) being used by the end user, are removed. In addition, if a certificate is required for access to a particular Link (identified in service data table 508), and the end user does not possess the required certificate, such Links are removed. The information relating to the country from which the end user has inputted the

code; the preferred language of the end user; the channel and/or channel type may be provided directly or automatically from the user's access media to site 604 (identified in Figure 6) either via a cookie stored on the access media, based on user selections that have previously been stored on the access media or other ways known in the art. U.S. Patent No. 6,314,457, which is hereby incorporated by reference, provides a description of exemplary access media that can be used in accordance with the present invention. The end user is then presented with the remaining Links. In some embodiments, companies may set priorities within the system (identified in service data table 508), thereby controlling the order in which Links are presented to end users.

[0034] A preferred embodiment of a system 600 that may be used to support the linking of unique identifiers to URLs using databases of the type described above is illustrated with reference to Figure 6. Variations on the components and configuration of the system 600 will be known to those skilled in the art and are within the scope of the present invention. The system described herein is highly scalable, thereby enabling it to handle a large number of Links. Further, the system is also designed to handle many end user lookups per second. The system also provides for both global and local load balancing, is stable and highly reliable.

[0035] System 600 includes a central relational database 601 (e.g., as described with reference to Figures 4A through 4C) for storing all information used in connection with system 600. Central relational database 601, for example, handles the companies' member accounts, which are used for online insertion or deletion of Links. Central relational database 601 is maintained on a central site 604 and one or more cache databases 603 (e.g., as described with reference to Figure 5) are maintained on cache servers 602. In a preferred embodiment, system 600 includes a standby relational database 605 on standby site 606 that is a real time mirror of central relational database 601. In case of a crash or scheduled downtime on the central relational database 601, the standby relational database 605 will take over.

[0036] Updates to the system 600 (for example, the addition of new Links) are handled through the central relational database 601 and are propagated to the cache-server databases 603. In particular, when changes are made in central relational database 601, that information is extracted by the master cache server 607 (recording the change in its transaction log database) and broadcasted through net layer 608 to the different database sites, from which the data is retrieved by the cache servers 602 and stored in cache databases 603. The transaction log works as a persistent buffer such that cache servers 602 that are brought on line can obtain transactions missed while offline.

[0037] Lookups are handled from the Internet through the global load balancing layer (GLB) 609 to whichever web server 610 at the site has the shortest latency and load. The web server 610

connects to any of the cache servers 602 at the site and performs the lookup in the various tables (as described with reference to Figures 5). The result of the lookup is then transmitted back to the end user client 611 via the web servers 610. As lookups are being performed, statistics for each lookup request are gathered by the cache servers 602. That data is then sent at regular intervals back to the master cache server 607 that inserts the data into the central relational database 601.

[0038] In the preferred embodiment, operations in the system 600 can be performed in two modes: read only or read write. When in read write mode, the operations to the database are guarded through transactions, whereas, in read only mode, the system will bypass normal transaction handling. In this manner, lookup operations will not interfere with insert or delete operations, thereby allowing lookup operations to be performed without transaction locks, which greatly enhances the performance of the system.

[0039] Company accounts are managed through an administration site internal to system 600, which is also connected to the central relational database 601. Company clients 612 (i.e. local clients at companies sponsoring web pages to which the URLs point) can connect securely using a standard web browser through one more or HTTPS servers (i.e. web servers 610) at each site. The connection is load balanced just as the end-user lookup requests. An application program interface (API) provides a remote way of accessing central relational database 604. This allows for company clients 612 that are running locally at the company to connect to and change the contents of the central relational database 604 (for example, to access their accounts and manage their Links). The API requests are directed from the web servers 610 to one of cache servers 602, that transmits the request back to one of master servers 607, that connects to the central relational database 604 and performs the request. The result is transmitted back from the cache master server 607 to the cache server 602, to the web server 610, that transmits it back to the API and the remote company client 612.

[0040] Thus, for example, a plugin to a desktop publishing program allows a newspaper editor to retrieve and link up barcodes that can be used in magazines. The plugin allows the editor to login to his account in the system 600 and directly manage his Links from within the desktop publishing program. Each barcode can be retrieved through the tool as an image to be printed in the magazine.

[0041] As mentioned previously, a preferred embodiment of the system 600 uses global load balancing in order to meet demands in different parts of the world. This is also important for ensuring low latency and high reliability. The global load balancers 609 that the system 600 uses act like a global web server that routes requests through to the most suitable site. In particular, when a DNS request is made to a global load balanced system (such as system 600), load balancers 609

intercept the request and, depending on different parameters such as server health, site capacity, and round trip time, return the IP address of the site which is most suitable. At each site, local load balancing is used to increase performance and ensure site level reliability. In particular, the local load balancer distributes the load between local servers at a site to increase the lookup capacity at a particular site. Extending performance of the system 600 is, thus, simplified. In the event a need arises to add more capacity to a particular site, new hardware can be added to that site. In the event more capacity is needed in another part of the world, a new site can be added to the global system. Other load balancing schemes will be known to those skilled in the art.

[0042] Figure 7 is a flow chart illustrating a method for directing an end-user to a location on a network. The end-user is associated with at least one of the following: an end-user country selection, an end-user language selection, a service category, an access media type, and the presence of an access certificate. The service category may be a product- or service-specific category or a third party service provider category. The service category in effect for a session may be chosen by the end user. The access media can be one of at least two types (at least one of the types being a bar code scanner). In step 701, a database that includes a plurality of entries is provided. Each of the entries is associated with a machine-readable code and includes one or more different network addresses. The different network addresses are associated with different providers (i.e. companies). In some embodiments, the entries further include one or more services, each service comprising a collection of network addresses grouped in accordance with at least one shared characteristic. The shared characteristic may be an identified profession; a country of origin; a language; or an access media device type, for example.

[0043] In step 702, code information associated with one or more scans of machine readable codes, made by the end user with a remote scanning device, is received at a network site. In step 703, in response to the scan information, entries associated with the machine-readable codes scanned by the end user are retrieved from the database. At least some of the entries include information that is not relevant to at least one of the end-user country selection, the end-user language selection, the service category, the access media type, and the presence of an access certificate. In step 704, based on at least one of the end-user country selection, the end-user language selection, the service category, the access media type, and the presence of an access certificate, one or more network addresses are selected from the retrieved entries. In step 705, the end user is connected with the one or more network addresses selected in step 704. In some embodiments, in step 706, the network addresses are sorted based on a provider-indicated priority.

